



FERMILAB-Conf-81/40-EXP  
7320.516

USE OF THE ECL-CAMAC TRIGGER PROCESSOR SYSTEM  
FOR RECOIL MISSING MASS TRIGGERS AT THE TAGGED  
PHOTON SPECTROMETER AT FERMILAB\*

J. Martin, S. Bracker, and G. Hartner  
University of Toronto  
Toronto, Canada

and

J. Appel and T. Nash  
Fermi National Accelerator Laboratory  
Batavia, Illinois 60510 USA

May 1981

\*Presented at the Topical Conference on the Application of Micro-processors to High-Energy Physics Experiments, Cern, Geneva, Switzerland, May 4-6, 1981



Use of the ECL-CAMAC Trigger Processor System for  
Recoil Missing Mass Triggers at the  
Tagged Photon Spectrometer at Fermilab

J. Martin, S. Bracker, G. Hartner  
University of Toronto  
Toronto, Canada

and

J. Appel, T. Nash  
Fermi National Accelerator Laboratory  
Batavia, Illinois USA

Abstract

A trigger processor in operation since May 1980 at the Tagged Photon Spectrometer at Fermilab will be described. The processor, based on the Fermilab ECL-CAMAC system, allows fast selection of high mass diffractive events from the total hadronic cross section. Data from a recoil detector, consisting of 3 wire chambers and 4 layers of scintillator concentric about a 1.5 m liquid hydrogen target, is digitized and presented to the processor within 3  $\mu$ sec. From the chamber data are found the vertices and angles of all recoiling tracks. The energy and particle identification ( $\pi, p, e$ ) of each track is determined by a fit to the energy deposits in the scintillator. If there is a single recoiling proton from the most upstream vertex, the forward missing mass is calculated using the incident photon energy and the energy and angle of the proton. Neutral patterns in the scintillator from a  $\pi^0$  or neutron are also recognized. Final triggering decisions are based on the forward mass, number of tracks at the primary vertex, total number of tracks, number of neutrals, etc. Total processing time is typically less than 10  $\mu$ sec. The processor has proven extremely reliable and the complex software for calculating the loads for the many memory look-up units is easy to use.

1. Introduction

In this paper we discuss a trigger processor which has been in operation at the Tagged Photon Laboratory Spectrometer at Fermilab for just over one year. The spectrometer is a large system consisting of a recoil detector, two magnets, drift chambers, segmented Cerenkov counters, segmented electromagnetic and hadronic calorimeters and a final wall of muon counters.<sup>(1)</sup> It was designed with full acceptance for incident photons in the 70-150 GeV range, with particular emphasis on the detection of final states containing charmed particles. From the beginning it was recognized that some sort of flexible, fast and selective high level triggering capability would be necessary to fully exploit the potential of this spectrometer. This led to the development of a versatile ECL-CAMAC trigger processing system,<sup>(2)</sup> whose first application is the processor described here.

2. Why a Trigger Processor?

We have been taking data at an electron beam energy of 137 GeV (mean tagged photon energy  $\sim 100$  GeV). With our allocation of the main

ring proton intensity the hadronic event rate in the 1.5 m liquid hydrogen target is typically 1000-1500 per one second machine pulse. A fast trigger in NIM logic requiring calorimeter energy deposit greater than 0.4 times the photon energy detects all of this rate. In addition, a small fraction of the enormous rate of  $e^+e^-$  pairs (200 x the hadronic cross-section) leak into this trigger, increasing the fast trigger rate to  $\sim 2000$  per pulse. CAMAC read-in of a typical event takes 3 milliseconds. Clearly, the fast trigger must be cut from 2000 to  $\lesssim 100$  events per second for recording in order to keep the data taking live-time reasonable. The problem then is how to achieve this reduction in a way which enhances the charm content (in particular, charm content which can be analyzed!) of the recorded data.

The cleanest possible way to produce charm in a photon beam is diffractively: the final state contains a charmed and an anti-charmed particle, whose decay products lie within the forward spectrometer acceptance, and a proton recoiling to the side. The forward going mass can be calculated from the incident photon energy and the angle and energy of the recoiling proton. We therefore decided for our first data run to configure a processor to find a single recoiling track from the main event vertex, identify it as a proton, calculate its angle and energy and then the forward mass within several microseconds. The main high level trigger demands a recoil proton giving a forward mass in the range 2 to 11 GeV. This reduces the fast trigger to 7% of the total hadronic cross-section rate, thus nicely matching our data recording capability.

### 3. The ECL-CAMAC Processor System

A full review and specification of this system can be found elsewhere.<sup>(2,3,4)</sup> However, the salient features will be described very briefly here in order to provide a basis for the discussion of our particular processor configuration in Section 5 below.

The design philosophy was to provide a flexible, modular unclocked system based on high speed ECL and large table look-up memories. The system is housed in modified CAMAC crates for easy communication with the experimental computer. Programming algorithms are accomplished by interconnecting modules at the front panels (as in NIM logic, but in this case with both data and control lines) and by preloading the table look-up memories. The path of the program is controlled by ready lines. Each module begins processing only when all its input ready lines are set, and sets its output ready as soon as its processing is finished.

Program flow thus proceeds as rapidly as module processing time (10-50 ns for most types of module) allows.

The two main work-horses of the system are the Memory Look-up Unit (MLU) and the Stack. The MLU contains up to 4096x16 bits of RAM. Front panel input bits are used to address the memory; the bit pattern at that address, which is some previously calculated, arbitrary function of the input bits, is presented at the front panel output data. The (input,output) field widths can be any one of (12,16), (13,8), (14,4), (15,2) or (16,1). A means of storing input data to the processor, or intermediate data from some internal loop for later use by some other part of the processor, is provided by the Stack. It contains 32 16-bit words of RAM which can be sequentially or randomly written and read. An internal priority scheme handles simultaneous read and write requests.

Other general purpose modules include a quad 4-bit function module, normally used as a 40 MHz scaler, ECL to NIM level converters to communicate with the experimental fast logic, a 3-fold Fanout unit, a double index Do Loop Indexer and a 3-fold General Logic Module (GLM), which can perform any one of the following operations on control signals:  $(A \div B) \div C \cdot G = \text{output}$ . Most modules can be controlled or have their memory and front panel signals examined by CAMAC. Single stepping through a processor program is thus possible; this makes debugging from the host computer relatively easy.

#### 4. The Recoil Detector and Associated Electronics

A drawing of the recoil detector is shown in Fig. 1. The inner part of the detector consists of three nesting, low mass (paper honeycomb-mylar sandwich construction), cylindrical proportional chambers with both anode ( $\phi$ ) and cathode ( $\theta$ ) readout. There are 656 cathode strips in hoops perpendicular to the beam along the 2 m length of each chamber. Cathode hits are latched and then scanned from upstream to downstream by fast digital processors, which find and calculate the width and centroid, to 1/2 wire spacing, of each cluster (typically 6 strips wide). This data from the three chambers is multiplexed and transmitted over a long cable to a receiver in the trigger processor at an average rate of 300 ns/cluster. The polar angle  $\theta$  of recoil tracks can be found from this cathode information with a resolution better than 6 mr.

Outside the chambers are four layers of scintillator: the inner two, A and B, are plastic NE110; the outer two, C and D, are liquid NE235A. Protons with kinetic energy less than 45 MeV are ranged out by

the liquid hydrogen and chamber material; thus the  $t$  acceptance begins at  $-0.09$ . The scintillator is divided into 15 azimuthal sectors so that events with several recoil tracks can be properly analysed. Light is read out at the upstream end only by 180 phototubes glued directly to the scintillator. On each end of the A layer counters is glued a fast timing tube as well. The end to end timing difference (EET) between these tubes gives the  $Z$  position of the particle entering each sector. These tube outputs are also used in conjunction with the tagged photon time to give the time of flight (TOF) from the target. Dynode signals from the light collecting tubes, EET and TOF are all digitized in  $1\ \mu\text{sec}$  by a system of fast ADCs and TDCs<sup>(5)</sup> and sent to the trigger processor.

The layer thicknesses were chosen to optimize energy resolution,  $\pi$ -p separation and neutral ( $n$  or  $\gamma$  from  $\pi^0$ ) recognition. For stopping protons, the energy resolution is always less than 5%. However, the chance that nuclear scattering will destroy the energy measurement is high. The A and B layers are therefore relatively thin. The kinetic energy of protons up to 650 MeV can be calculated to better than 10% from the energy deposits in these two layers alone. Pions can be cleanly separated from protons up to 600 MeV using A,B,C and D energy deposits, and up to 300 MeV using A and B only. Neutrals are identified by light appearing in layers beyond A, but not in A.

The trigger processor expects the calibration for all 15 sectors to be the same. This has proven easy to maintain, as phototube stability is quite good. Calibration is carried out offline every few weeks using a large sample of recoil protons from normal running data. Gain and attenuation parameters are fitted for each sector, and any phototube voltage adjustments necessary to bring these back to the average values are calculated and made. Typically, gains are stable within 5% over many weeks.

The resolution in angle and energy of the recoil proton, discussed above, and in incident photon energy ( $\sim 4\%$ ) leads to a forward mass resolution better than 300 MeV for masses greater than 2 GeV.

## 5. The Recoil Processor

Superimposed on the side view of the recoil detector in Fig. 1 is a typical sort of event which the processor must analyse. A recoil proton emerges at small  $\theta$  from the primary vertex. One of the forward secondaries has an interaction with a 3 charged particle recoiling system (e.g.,  $p\pi^+\pi^-$ ). A  $\delta$  electron from another secondary multiple scatters out through PWC I and J. Not shown are light deposits from a

$\pi^0$  or neutron if the primary vertex is  $p\pi^0$  or  $n\pi^+$  instead of  $p$ . Background electrons scattered into the detector from low energy pair or Compton production from the huge flux of low energy bremsstrahlung and synchrotron radiation photons can also occur within the time resolution of the main event gate. Of course, many clean events with only a single recoiling proton also occur, but the processor must be able to handle the complex general case.

A much simplified schematic of the processor configuration is shown in Fig. 2. There are several loops operating simultaneously. The track finder loop finds all possible track combinations from the PWC cathode data and stores the vertex position ( $V_z$ ) and angle of each track in two stacks. The track loop can begin as soon as at least one centroid from each chamber has arrived ( $2-3 \mu\text{sec}$ ). The entry of the first track into the track stacks signals the start of the main loop of the processor, which matches the track to a scintillator sector and determines its energy and species ( $\pi, p, e$ ). This loop makes use of ADC, EET and TOF information from the fast ADC-TDC system, which has been loaded into stacks within  $3 \mu\text{sec}$ . The sector matching itself is a loop within the main loop, in which the  $Z$  position at the A layer from the track projection and from EET are compared for each sector in turn until a match is found. This provides the  $\theta$ - $\phi$  correlation between the chamber tracks ( $\theta$  only) and the scintillator energy deposits ( $\phi$  only). Outside the main loop, another loop cycles through the ADC information from all 15 sectors looking for neutral patterns. After all tracks have been found and processed by the main loop, the scaler information, neutral information and missing mass (calculated from  $\theta$  and  $E$  of the most upstream proton data and the tagged photon energy) is presented to a final MLU for a triggering decision.

#### Some More Detailed Sub-Section Descriptions

Some details of the innards of the processor will now be given in order to illustrate the complexity of the analysis which it performs. However, the editor's patience will not be taxed with a cable by cable, module by module description.

i) The Track Finder : The track finding subsystem is described completely in references 2,3 and 4. The system calculates the line projection on PWC J for each combination of PWC I and K centroids. A match with hits in a hit array generated from PWC J centroids (allowance for chamber alignment is made) indicates a track, and its slope and vertex are sent to the track stacks. The loop time is 130 ns, so all tracks are found in  $130 \times (\text{no. of I centroids}) \times (\text{no. of J centroids})$ .

The PWC data receiver has a PROM which allows wide clusters due to overlapping tracks to be split into two. Because of the large variation in cluster width for single tracks of widely varying angle and ionization loss, this feature is not used.

ii) Sector Finder Loop : This loop, shown in Fig. 3, illustrates succinctly the simple use of 3 standard modules to solve quickly a relatively complicated problem ( $\theta$ - $\phi$  track correlation). At the beginning of each main track loop, the EET stack sets its output ready (OR) and output data (OD) from the first word, which contains sector no. M and its associated time  $T_M$ . The stack has been loaded previously only with data from sectors with physical times by a sparsifying circuit in the EET receiver. Thus time is not wasted looping over sectors with no tracks. The sector finder MLU waits for  $Z_A$  (the projected position of the track in layer A, calculated in an earlier MLU in the track loop) to be ready, then does a  $T_M, Z_A$  comparison. If a match is found within the EET resolution, the MLU OD will indicate YES, and the track loop proceeds with M and M ready to pick up the ADC information from sector M. If NO, the stack receives a sequential read request (SRR) which toggles OR (IR to the MLU) and presents its next word on OD. The loop thus continues; the loop time is  $\sim 80$  ns. Note that all timing is controlled by the readies. The only, and trivial, restriction on cable lengths is that a ready cable must be at least as long as its accompanying data. When all available  $T_M$ s have been examined, data exhausted (DE) from the stack is set. The track loop then proceeds by a NO MATCH branch.

iii) Energy Determination and Particle Identification : Figure 4 shows the data flow through this section of the track loop. Readies are omitted for clarity. Note that data can be used in more than one module by simple daisy-chaining of buffered inputs ( $\cos\theta$  and  $E_1$ ).

We first discuss the  $E_1$  MLU as an example of the careful considerations required to make optimum use of MLUs. The  $E_1$  MLU does a fit to the energy deposits in A and B ( $E_A$  and  $E_B$ ) from matching sector M, for both  $\pi$  and p hypotheses. The best kinetic energy ( $E_1$ ) from the proton fit is output, along with 2 bits,  $\pi_1$  and  $p_1$ , which indicate whether or not the fits passed  $\pi$  and p  $\chi^2$  tests. It is possible that both of these bits are on, or off, because of resolution.

Resolution here is the key word: it governs the number of bits needed for each variable, and, most importantly, the manner in which the variable data is packed into the bits. We need at least 5 bit resolution in kinetic energy (5%) for the missing mass calculation (see Section 4). Thus the output field,  $E_1, \pi_1, p_1$ , is 7 bits wide, which

implies that only 13 input bits can be used (see Section 3). The optimum split is 5 bits each for  $E_A$  and  $E_B$  and 3 bits for  $\csc\theta$ , which is proportional to the track path length through the scintillator block and on which the fit depends.  $E_A$  and  $E_B$  are previously calculated in MLUs from the light pulse heights (7 bits, linear) using the average gain and attenuation parameters of the 15 sectors. Packing is in the form  $E_A = c_1 \exp(\text{bin no.})^{1/2}/c_2$ , where  $c_1$  and  $c_2$  are appropriate constants to compress the dynamic range into the 32 bins (bin no. = 0,1,...,31). This very non-linear packing allows the digital resolution (1/2 bit) to match the energy resolution of the detector throughout the range. The  $\csc\theta$  packing is  $\exp c_3(\text{bin no.})$ , with  $c_3$  chosen to provide binning roughly linear in path length between  $20^\circ$  and  $90^\circ$ .

To produce the table look-up in the  $E_1$  MLU required calculation of the expected energy deposits for protons and pions as a function of incident angle and kinetic energy. (Energy deposit in the scintillator follows the Bethe-Bloch equation; this was checked with a detector prototype in a low energy proton beam.) Complex fitting procedures were written to parameterize these results and provide the inverse fit, kinetic energy  $E_1 = f(E_A, E_B, \csc\theta)$ . A much more complicated algorithm than that shown above for  $E_A$  was necessary to pack  $E_1$  into 5 bits to match the kinetic energy resolution, itself a complicated function of the resolution in  $E_A$  and  $E_B$ . The physics calculation done in the  $E_1$  MLU is thus non-trivial. The enormous power of the MLU can be appreciated by noting that the PDP 11-55 programming to produce the  $E_1$  MLU load took many hours of CPU time; the on-line look-up evaluation of  $E_1$  takes only 50 ns.

A similar calculation of kinetic energy is performed in the  $E_2$  MLU (not shown in Fig. 4) from the C and D energy deposits. The E MLU then provides consistency checking and a combined fit with  $E_1$  and  $E_2$  for the best estimate of kinetic energy,  $E$ , for the proton hypothesis.

Particle identification comes from the combined output of the TOF and P MLUs. TOF MLU makes a decision based on time of flight only for those particles which stop in layer A (low  $E_1$ ). The TOF resolution ( $\sim 0.8$  ns) is such that clean distinction between p or electron is often not possible. Both output bits are then set to zero. The P MLU functions as a simple combinatorial logic circuit. It decides from the  $\pi_1 p_1 \pi_2 p_2$  bits to assign the track to one of three output categories:  $P_{\text{poss}}$ , possibly a proton (e.g., 0011);  $P_{\text{def}}$ , definitely a proton (e.g., 0111); or INCON, inconsistent with either p or  $\pi$  (e.g., 0001). If none of the 3 output bits is set, the track is considered to be a pion.

An example of the flexibility of the system is provided by the



INCON circuit. This was added at a late stage when the EET resolution was found to be worse than expected (~11 cm in Z position instead of 5). There is a reasonable probability that more than one sector will match a PWC track projection within this somewhat degraded resolution. The first match found may not be the right one, and INCON=1 will likely result. This signal returns to the sector finder loop to seek another match (see Fig. 3). Note that this is equivalent to branching from outside to inside a FORTRAN DO loop, illegal in FORTRAN, but not here, because the EET stack has remembered the loop index!

Figure 4 illustrates one final point, that timing considerations are sometimes necessary. The OR of the P MLU is used as the input ready to the vertex subsection. It must be delayed to wait for E (minimal, since E, TOF and P MLUs work in parallel) and to wait for possible blocking by INCON. There are several other similar, some more complex, situations where branching occurs in the main loop. Such timing considerations present no problem to any experimental physicist experienced in the use of NIM logic.

iv) Vertex Selection and Scaling : The vertex subsection consists of a special purpose vertex parameter module (VPM) and an MLU. The VPM stores  $\cos\theta$  and target vertex position  $V_z$  of the first track through the main loop. For each subsequent track the difference between its vertex and the stored vertex is calculated. The vertex MLU uses this to determine whether the current track comes from a different vertex than that of the stored track. It issues a STORE to the VPM if the new track is at a new, more upstream, vertex (such a situation occurs in the event shown in Fig. 1) or if it is a proton at the same vertex. In this way  $\cos\theta$  and  $V_z$  of the most upstream proton are remembered for later use in the missing mass MLU. Particle identity bits pass through the vertex MLU and are combined with the vertex information to provide output data to be scaled.

The scalers include the following:

- a)  $P_{\text{poss}}$ ,  $P_{\text{def}}$ , and no. of tracks, all at upstream vertex only.
- b) Total protons, total pions, total tracks and no. of new vertices.
- c) NO MATCHs, backward tracks, neutrals.

Scalers in group a) must be cleared whenever a new upstream vertex is found. This CLEAR is part of the vertex MLU OD.

v) Final Triggering : When all tracks have been processed, control is passed to the trigger MLU. Input data includes 12 bits of selected scaler information and 2 bits from the missing mass MLU indicating four possible forward mass regions. Fourteen input bits allow 4 output bits, which are used for four different triggers. Two of these are the main

diffractive triggers. They demand a single track ( $P_{\text{poss}}$ ) at the upstream vertex, no neutrals (to eliminate  $n\pi^+$  and  $p\pi^0$  recoils), no more than 2 NO MATCH and backward tracks and total number of tracks less than 8. The mass ranges for these two triggers are 2.0 to 5.5 GeV and 5.5 to 11.0 GeV. A third trigger, which is prescaled by a large factor, accumulates data in the 0-25 GeV range to aid in off-line acceptance calculations. The fourth trigger, also somewhat prescaled, simply demands  $\geq 3$  tracks at the primary vertex. This is an attempt to record some associated charm production events. The trigger has reasonable acceptance for such events, since the recoil system contains a charmed baryon which can decay with high multiplicity.

Total event trigger processing time is a complicated overlap function of the data read-in time ( $\sim 3 \mu\text{sec}$ ), the trackfinding time  $(\text{no. of tracks})^2 \times 130 \text{ ns}$  and the main track loop time ( $1.2 \mu\text{sec}$ ). In general this is less than  $10 \mu\text{sec}$ . There is an abort time which cuts off processing at  $35 \mu\text{sec}$  for the occasional noisy events in which the PWCs are blitzed with hits.

## 6. Memory Look-up Unit Software

The recoil trigger processor uses 19 MLUs containing approximately 64K 16-bit words. Generating  $2^{32}$  bits of appropriate information so the processor can function is a daunting prospect. Fortunately for the physicist user, a software package to make this task as painless as possible is part of the ECL-CAMAC system.

The main job of the user is to provide a parameter file and a FORTRAN subroutine for each MLU. The parameter file holds constants (150 in this application) used in building the MLU loads. These constants, some used by more than one MLU, are frequently updated during experimental set-up and testing, and are gathered together in one file for each reference. A skeleton form of the MLU FORTRAN is part of the software package. The user need only flesh it out with the following information: the MLU name, the substructure and names of its input and output fields, the names of parameters and the code (physics, logical and/or geometric, etc.) required to generate the output bit pattern from the input field.

A routine of the software package uses the parameter file and MLU subroutines to build disk files of the output fields for every possible input field pattern. Another general utility loads the MLUs from the disk files via CAMAC, and reads back the load for verification.

The system software package also has convenient bookkeeping

procedures to quell the general tendency toward chaos and anarchy in software revision. The current state of the processor load can thus be readily understood.

## 7. Debugging and Results

Debugging proceeded in two steps, the first a thorough electronics check-out, the second a physics analysis with real data.

The first stage was accomplished with simulated events. PWC, EET and TOF data were loaded into the processor via test modules, the processor was triggered, and final results of the track stacks and scalers were compared with expectation. A special test load of the TOF MLU was used to identify particle type based on the input TOF bits alone. ADC information was ignored by loading the P MLU to give zero output for all input combinations. Particularly pathological events were invented with multiple vertices and multiple crossing tracks of all particle types at each vertex. These tested every possible branch within the main loop, the sector finder loop and the vertex-scaler subsystem. The track finder was exhaustively tested with tens of thousands of randomly generated tracks. The failure rate was less than  $3 \times 10^{-5}$ .

It should be noted that this stage of debugging was relatively easy owing to the use of the extensible language system FORTH. This system provides in-line compilation and execution of commands, which can be composed as the need arises, and in an extremely flexible manner, from already defined commands in the system. Thus, as debugging continued, a powerful hierarchic structure of commands was built up, from the very low level (e.g., read X and Q response from some ECL-CAMAC module) to very high level (e.g., one command to load a complicated event pattern, trigger the processor and type the scaler results).

Events which failed could be followed through the processor in steps by inhibiting processing at various points in the track loop and examining the state of front panel indicating lights, and reading input and output addresses to MLUs and stack contents with FORTH. Such events could also be triggered repeatedly at a convenient rate, say 1 kilohertz, for oscilloscope monitoring of ready and data lines. Most of the problems found in this way were traced to cabling errors or incorrect MLU loads. Hardware failure of the modules has been rare. One MLU was replaced just before the 6-month run now near completion; since then, no other module has died.

The physics debugging involved taking data with the processor not

in the trigger, but with the results of its analysis read out and recorded. Diagnostic stacks were used at various points in the main loop to accumulate for read-out such quantities as  $E_A$ ,  $E_1$ ,  $E_2$ ,  $E$ ,  $\cos\theta$ , etc. for all tracks in each event. Histograms of these quantities and correlations of the scaled quantities for thousands of events were compared with off-line reconstruction of the recoil data. This comparison revealed several programming bugs in the MLU loading software. Many such histograms could be shown, but perhaps the most dramatic is seen in Fig. 5. Here is plotted the off-line determination of the forward missing mass,  $M_x$ , for events from the two different processor triggered mass ranges at a beam energy of 170 GeV. It can be seen that the  $M_x$  resolution, which depends on all parts of the recoil detector and processor working correctly, is very good. The distribution falls off at larger mass because of acceptance.

## 8. Conclusions

The recoil trigger processor has been a most successful application of the ECL-CAMAC trigger processor system. More than 15 million high energy photoproduction events triggered with the processor have so far been recorded. The charmed particle physics which these events hold is expected to be rich.

It is our pleasure to acknowledge the superb effort of the engineering staff in the Beam Systems Group at Fermilab, in particular E. Barsotti, M. Haldeman, R. Hance, W. Haynes, T. Soszynski, and K. Treptow, in the design and building stages of the ECL-CAMAC system. Their support during the implementation and debugging of the recoil processor was invaluable. We also thank our colleagues on Experiment 516 for their good-humored, skeptical at first, encouragement as the system was brought on line.

## References

1. J. Appel, D. Bartlett, S. Bracker, G. Hartner, G. Kalbfleisch, G. Lüste, P. Mantsch, J. Martin, R. Morrison, T. Nash, U. Nauenberg, D. Ritchie, K. Stanfield, S. Yellin, The Tagged Photon Magnetic Spectrometer: Facility Design Report, Fermilab (1977, unpublished). The spectrometer was built and is being run by a collaboration of physicists from the Universities of California at Santa Barbara, Carleton, Colorado, Oklahoma, Toronto and from Fermilab and the National Research Council of Canada Laboratory.
2. E. Barsotti, J. Appel, S. Bracker, M. Haldeman, R. Hance, B. Haynes, D. Kline, J. Krebs, J. Maenpaa, T. Nash, T. Soszynski, K. Treptow, ECL-CAMAC Trigger Processor Documentation, Fermilab TM-821 (1978).
3. E. Barsotti, J. Appel, S. Bracker, M. Haldeman, R. Hance, B. Haynes, J. Maenpaa, T. Nash, T. Soszynski, K. Treptow, A Modular Trigger Processing System for High Energy Physics Experiments, IEEE Transactions on Nuclear Science, Vol. NS-26, No. 1, 686 (1979).
4. T. Nash, Invited Review Talk, this conference.
5. Developed by C. Kerns, Fermilab.




Figure 1: Recoil Detector at the Tagged Photon Spectrometer.

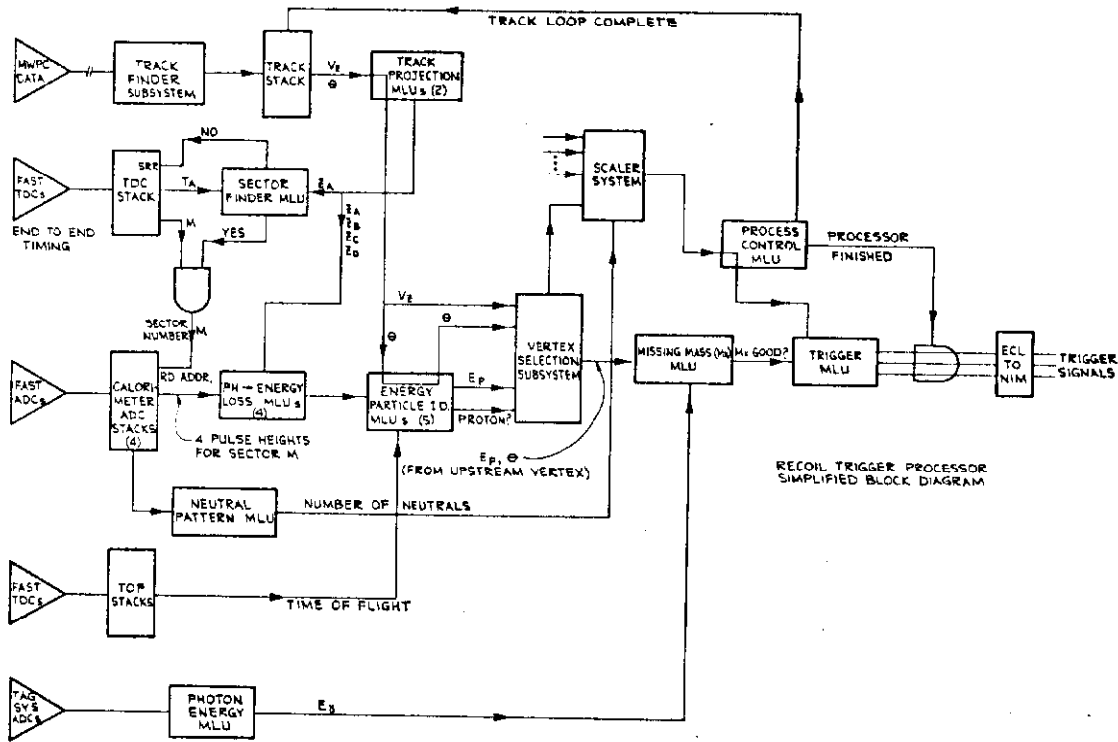



Figure 2: Recoil Trigger Processor Simplified Schematic.






Figure 5: Offline Reconstruction of Missing Mass from events in the two diffractive triggers of the Recoil Processor.